

An FPTAS for the Parametric Knapsack Problem[☆]

Michael Holzhauser^{a,*}, Sven O. Krumke^a

^aUniversity of Kaiserslautern, Department of Mathematics
Paul-Ehrlich-Str. 14, D-67663 Kaiserslautern, Germany

Abstract

In this paper, we investigate the parametric knapsack problem, in which the item profits are affine functions depending on a real-valued parameter. The aim is to provide a solution for all values of the parameter. It is well-known that any exact algorithm for the problem may need to output an exponential number of knapsack solutions.

We present a fully polynomial-time approximation scheme (FPTAS) for the problem that, for any desired precision $\varepsilon \in (0, 1)$, computes $(1 - \varepsilon)$ -approximate solutions for all values of the parameter. This is the first FPTAS for the parametric knapsack problem that does not require the slopes and intercepts of the affine functions to be non-negative but works for arbitrary integral values. Our FPTAS outputs $\mathcal{O}(\frac{n^2}{\varepsilon})$ knapsack solutions and runs in strongly polynomial-time $\mathcal{O}(\frac{n^4}{\varepsilon^2})$. Even for the special case of positive input data, this is the first FPTAS with a strongly polynomial running time. We also show that this time bound can be further improved to $\mathcal{O}(\frac{n^2}{\varepsilon} \cdot A(n, \varepsilon))$, where $A(n, \varepsilon)$ denotes the running time of any FPTAS for the traditional (non-parametric) knapsack problem.

Keywords: knapsack problems, parametric optimization, approximation algorithms

1. Introduction

The knapsack problem is one of the most fundamental combinatorial optimization problems: Given a set of n items with weights and profits and a knapsack capacity, the task is to choose a subset of the items with a maximum profit such that the weight of these items does not exceed the knapsack capacity. The problem is known to be weakly \mathcal{NP} -hard and solvable in pseudo-polynomial time. Moreover, several constant factor approximation algorithms and approximation schemes have been developed for the problem [1–5] (cf. [6] for an overview).

In this paper, we investigate a generalization of the problem in which the profits are no longer constant but affine functions depending on a parameter $\lambda \in \mathbb{R}$. More precisely, for a knapsack with capacity W and for each item i in the item set $\{1, \dots, n\}$

with *weight* $w_i \in \mathbb{N}_{>0}$, the *profit* p_i is now of the form $p_i(\lambda) := a_i + \lambda \cdot b_i$ with $a_i, b_i \in \mathbb{Z}$. The resulting optimization problem can be stated as follows:

$$\begin{aligned} p^*(\lambda) = \max \quad & \sum_{i=1}^n (a_i + \lambda \cdot b_i) \cdot x_i \\ & \sum_{i=1}^n w_i \cdot x_i \leq W \\ & x_i \in \{0, 1\} \quad \forall i \in \{1, \dots, n\} \end{aligned}$$

The aim of this *parametric knapsack problem* is to return a partition of the real line into intervals $(-\infty, \lambda_1], [\lambda_1, \lambda_2], \dots, [\lambda_{k-1}, \lambda_k], [\lambda_k, +\infty)$ together with a solution x^* for each interval such that this solution is optimal for all values of λ in the interval. The function mapping each $\lambda \in \mathbb{R}$ to the profit of the corresponding optimal solution is called the *optimal profit function* and will be denoted by $p^*(\lambda)$ in the following. It is easy to see that p^* is continuous, convex, and piecewise linear with breakpoints $\lambda_1, \dots, \lambda_k$ [6].

Clearly, since the parametric knapsack problem is a generalization of the traditional (non-parametric)

*Corresponding author. Fax: +49 (631) 205-4737. Phone: +49 (631) 205-2511

Email addresses: holzhauser@mathematik.uni-kl.de (Michael Holzhauser), krumke@mathematik.uni-kl.de (Sven O. Krumke)

knapsack problem, it is at least as hard to solve as the knapsack problem. In fact, it was shown that, even in the case of integral input data, the minimum number of breakpoints of the optimal profit function can be exponentially large, so any exact algorithm may need to return an exponential number of knapsack solutions [7]. In this paper, we are interested in a *fully polynomial time approximation scheme* for the parametric knapsack problem. We will show that, for any desired precision $\varepsilon \in (0, 1)$, a polynomial number of intervals suffices in order to be able to provide a $(1 - \varepsilon)$ -approximate solution for each $\lambda \in \mathbb{R}$.

Without loss of generality, we may assume that $w_i \leq W$ for each $i \in \{1, \dots, n\}$ since otherwise we are not able to choose item i at all. However, note that we do not set any further restrictions on the profits, i.e., the parameters a_i and b_i . In particular, each profit may become negative for some specific value of λ . It is even possible that there are no profitable items at all for some values of λ .

1.1. Previous work

A large number of publications investigated parametric versions of well-known problems. This includes the parametric shortest path problem [8–11], the parametric minimum spanning tree problem [12, 13], the parametric maximum flow problem [14–16], and the parametric minimum cost flow problem [7] (cf. [17] for an overview). The parametric knapsack problem considered here was first investigated by Carstensen [7]. She showed that the number of breakpoints of the optimal profit function can become exponentially large in general. If the parameters are integral, the number of breakpoints can still attain a pseudo-polynomial size. The first specialized exact algorithm for the problem was presented by Eben-Chaime [18], who showed that the problem can be solved in $\mathcal{O}(knW)$ time, where k denotes the number of breakpoints of the optimal profit function p^* .

The first approximation scheme for the problem was recently published by Giudici et al. [17]. The authors presented a generalization of the standard polynomial-time approximation scheme for the knapsack problem, resulting in a PTAS for the problem with a running time in $\mathcal{O}(\frac{1}{\varepsilon^2} \cdot n^{\frac{1}{\varepsilon}+2})$. In the special case of positive values of λ as well as non-negative values of a_i and b_i for each $i \in \{1, \dots, n\}$, the authors show that an algorithm of Erlebach

et al. [19] for the bicriteria knapsack problem can be used to obtain an FPTAS for the parametric knapsack problem running in $\mathcal{O}(\frac{n^3}{\varepsilon^2} \cdot \log^2 \text{UB}_{\max})$ time, where UB_{\max} denotes an upper bound on the maximum possible profit with respect to both of the profit functions $\sum_{i=1}^n a_i \cdot x_i$ and $\sum_{i=1}^n b_i \cdot x_i$.

1.2. Our contribution

We present the first FPTAS for the parametric knapsack problem without the restriction to non-negative input data. In particular, we show that we only need a total number of $\mathcal{O}(\frac{n^2}{\varepsilon})$ intervals to approximate the problem (which, itself, may need an exponential number of intervals as described above) and that we can compute an approximate solution for each interval in $\mathcal{O}(\frac{n^2}{\varepsilon})$ time, yielding an FPTAS with a strongly polynomial running time of $\mathcal{O}(\frac{n^4}{\varepsilon^2})$. Our algorithm is the first FPTAS for the problem with a strongly polynomial running time, being superior to the PTAS of Giudici et al. [17] for $\varepsilon \leq 0.5$ and, in the special case of positive input data, superior to their FPTAS for large input values. In a second step, we improve this result to a running time of $\mathcal{O}(\frac{n^2}{\varepsilon} \cdot A(n, \varepsilon))$, where $A(n, \varepsilon)$ denotes the running time of any FPTAS for the traditional knapsack problem. Using the FPTAS of Kellerer and Pferschy [4, 5], this yields a time bound of $\mathcal{O}(\frac{n^3}{\varepsilon} \log \frac{1}{\varepsilon} + \frac{n^2}{\varepsilon^4} \log^2 \frac{1}{\varepsilon})$ for the parametric knapsack problem.

1.3. Organization

The results of this paper are divided into three main parts. In Section 2, we show how we can generalize the well-known greedy-like $\frac{1}{2}$ -approximation algorithm for the traditional knapsack problem to the parametric setting and how the resulting profit function can be “smoothened” such that it becomes convex and continuous without losing the approximation guarantee. This will be the key ingredient for the parametric FPTAS, which will be presented in Section 3. We will first recapitulate a basic FPTAS for the traditional knapsack problem in Section 3.1 and then extend it to the parametric case in Section 3.2, presenting a first time bound for the resulting FPTAS. In Section 3.3, as a main result of the paper, we present an improved analysis yielding the claimed running time of the FPTAS. Finally, in Section 4, we show that it suffices to solve the corresponding subproblems only approximately so that we can incor-

porate traditional FPTASs, which improves the running time of our algorithm to the claimed one.

2. Obtaining a parametric $\frac{1}{2}$ -approximation

The parametric FPTAS will rely on a convex and continuous $\frac{1}{2}$ -approximation of the optimal profit function $p^*(\lambda)$, i.e., a parametric $\frac{1}{2}$ -approximation algorithm for the parametric knapsack problem. We will therefore present such an algorithm in this section and describe how we can guarantee these properties of the function.

2.1. Traditional $\frac{1}{2}$ -approximation algorithm

The basic (non-parametric) $\frac{1}{2}$ -approximation algorithm proceeds as follows: In a first step, the algorithm sorts the items in decreasing order of their ratios $\frac{p_i}{w_i}$, which can be done in $\mathcal{O}(n \log n)$ time. It then packs the items in this ordering until the next item k with $k \geq 2$ would violate the knapsack capacity (or until there are no items left), yielding a feasible solution x' . The algorithm either returns x' or, if better, the solution containing only an item with the largest profit $p^{(\max)}$. If x^* denotes an optimal solution to the given knapsack instance, x^A the solution returned by the above algorithm, and x^{LP} a solution to the LP-relaxation of the problem, we get that

$$\begin{aligned} p^A := p(x^A) &:= \sum_{i=1}^n p_i \cdot x_i^A = \max \left\{ p^{(\max)}, \sum_{i=1}^{k-1} p_i \right\} \\ &\geq \frac{1}{2} \cdot \sum_{i=1}^k p_i \geq \frac{1}{2} \cdot p(x^{LP}) \geq \frac{1}{2} \cdot p(x^*), \end{aligned}$$

so x^A is a $\frac{1}{2}$ -approximation. We refer to [6] for further details on this standard algorithm.

2.2. Parametric $\frac{1}{2}$ -approximation algorithm

In the parametric knapsack problem, the profits are affine functions of the form $p_i(\lambda) = a_i + \lambda \cdot b_i$ such that the optimal profit p^* changes with λ . However, note that the solution x' of the traditional $\frac{1}{2}$ -approximation algorithm only depends on the ordering of the items and, thus, remains constant as long as this ordering does not change. Moreover, two items can only change their relative ordering if their profit functions intersect, yielding $\mathcal{O}(n^2)$ intervals I'_j , within which the ordering of the items remains unchanged. For all values of λ in such an

interval I'_j , the algorithm computes the same solution x' , which has a profit of the form $p^{(j)}(\lambda) := \alpha^{(j)} + \lambda \cdot \beta^{(j)}$. For each $\lambda \in I'_j$, the $\frac{1}{2}$ -approximation algorithm either returns x' with profit $p^{(j)}(\lambda)$ or the most valuable item only. One possibility to obtain a parametric $\frac{1}{2}$ -approximation algorithm would be to consider each interval I'_j separately and to divide it into subintervals, depending on whether $p^{(j)}(\lambda)$ or $p^{(\max)}(\lambda)$ is larger, where $p^{(\max)}$ denotes the profit of the most valuable item (which, again, now depends on λ). However, the resulting piecewise linear function $p^A(\lambda)$ is not necessarily continuous or convex, which will be required later, though (see Figure 1).

Instead, we ignore the intervals I'_j and only consider the above affine functions $p^{(j)}(\lambda) = \alpha^{(j)} + \lambda \cdot \beta^{(j)}$. Let S denote the set of all such functions together with the function $p^{(0)}(\lambda) := 0$ and each profit function $p_i(\lambda)$. By computing the upper envelope of the $\mathcal{O}(n^2)$ functions in S , we obtain a function φ , which is based on feasible solutions whose profit is not smaller than $p^A(\lambda)$ at each $\lambda \in \mathbb{R}$ (see the dotted curve in Figure 1). By standard arguments, it follows that φ is convex, piecewise linear, and continuous as it is the pointwise maximum of affine functions. For m functions, the upper envelope can be computed in $\mathcal{O}(m \log m)$ time as shown¹ by Hershberger [20]. Within the same time bound, we can sort the resulting intervals by increasing values of their left boundary. Hence, we obtain a piecewise linear, continuous, and convex $\frac{1}{2}$ -approximation φ with $\mathcal{O}(n^2)$ breakpoints in $\mathcal{O}(n^2 \log n)$ time. In the following, we will refer to the intervals between the breakpoints of φ as I_1, \dots, I_q .

3. Obtaining a parametric FPTAS

Before we explain the parametric FPTAS in detail, we first recapitulate the basic FPTAS for the traditional (non-parametric) knapsack problem as introduced by Lawler [21] since its way of proceeding is crucial for the understanding of the parametric version.

¹Strictly speaking, the author proves the result for finite line segments and not for straight lines. However, it is easy to compute upper and lower bounds for the smallest and largest possible intersection point of two of the involved functions, respectively, and to reduce the problem to the resulting interval.

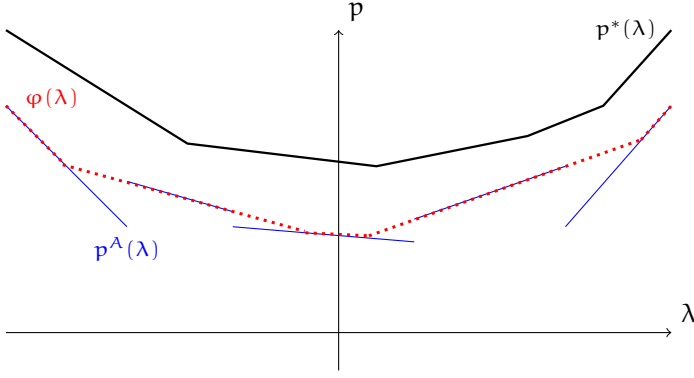


Figure 1: The optimal solution value p^* (black thick line), the profit function p^A of the $\frac{1}{2}$ -approximation algorithm (blue straight lines), and the “smoothed” curve φ (red dotted lines).

3.1. Traditional FPTAS

Consider the case of some fixed value for λ , such that the profits have a constant, but possibly negative value p_i . The basic FPTAS for the traditional knapsack problem is based on a well-known dynamic programming scheme, which was originally designed to solve the problem exactly in pseudo-polynomial time: Let P denote an upper bound on the maximum profit of a solution to the given instance. For $k \in \{0, \dots, n\}$ and $p \in \{0, \dots, P\}$, let $w(k, p)$ denote the minimum weight that is necessary in order to obtain a profit of exactly p with those items in the item set $\{1, \dots, k\}$ that have a non-negative² profit. For $k = 0$, we set $w(0, p) = 0$ for $p = 0$ and $w(0, p) = W + 1$ for $p > 0$. For $k \in \{1, \dots, n\}$ and for the case that $p_k \in \{0, \dots, p\}$, we compute the values $w(k, p)$ recursively by $w(k, p) = \min\{w(k-1, p), w(k-1, p - p_k) + w_k\}$, representing the choice to either not pack the item or to pack it, respectively. Else, if $p_k \notin \{0, \dots, p\}$, we set $w(k, p) = w(k-1, p)$ since we can omit negative item- and knapsack-profits. The largest value of p such that $w(n, p) \leq W$ then yields the optimal solution to the problem. The procedure runs in pseudo-polynomial time $\mathcal{O}(nP)$.

The idea of the basic FPTAS is to scale down the item profits p_i to new values $\tilde{p}_i := \lfloor \frac{p_i}{M} \rfloor$, where $M := \frac{\varepsilon \cdot \bar{p}}{n}$ for some value \bar{p} fulfilling $\frac{1}{2} \cdot p^* \leq \bar{p} \leq p^*$. Instead of setting $\bar{p} := p^A$ as it is done in the traditional FPTAS, we can alternatively use our improved 2-approximate solution φ . Since the maximum pos-

sible profit \tilde{P} is then given by

$$\tilde{P} \leq \sum_{i=1}^n \tilde{p}_i \leq \sum_{i=1}^n \frac{n \cdot p_i}{\varepsilon \cdot \varphi} \leq \frac{n}{\varepsilon} \cdot \frac{p^*}{\varphi} \leq \frac{2n}{\varepsilon},$$

the procedure runs in polynomial time $\mathcal{O}(\frac{n^2}{\varepsilon})$. The crucial observation is that we only lose a factor of $(1 - \varepsilon)$ by scaling down the profits, so the solution obtained by the above dynamic programming scheme applied to the scaled profits yields a $(1 - \varepsilon)$ -approximate solution for the problem (see [6, 21] for further details on the algorithm).

3.2. Parametric scaling

Although the parametric FPTAS is based on the basic FPTAS, the instance parameters now depend on λ and, thus, change while λ increases. In particular, both the item profits and φ now depend on λ , so the scaled profits $\tilde{p}_i(\lambda) := \lfloor \frac{n \cdot p_i(\lambda)}{\varepsilon \cdot \varphi(\lambda)} \rfloor$ have a highly non-linear behavior. Nevertheless, similar to the parametric $\frac{1}{2}$ -approximation considered in Section 2.2, the solution returned by the FPTAS does not change as long as the profit $\tilde{p}_i(\lambda)$ of each item remains constant. Hence, if I'_j denotes an interval such that the scaled profits \tilde{p}_i remain constant for each $\lambda \in I'_j$, we can evaluate the dynamic programming scheme with the profits \tilde{p}_i to obtain a $(1 - \varepsilon)$ -approximate solution for the interval I'_j . The proof of the polynomial running time and the approximation guarantee remain unchanged.

It remains to show how we can divide the real line into a polynomial number of intervals such that the profits remain constant in each interval. The basic FPTAS will then be evaluated for each such interval subsequently (using the corresponding constant scaled profits) in order to obtain $(1 - \varepsilon)$ -approximate solutions for the whole real line.

One natural idea would be to build on those intervals described in Section 2.2 for which $\varphi(\lambda)$ behaves like an affine function: Let I_1, \dots, I_q with $q \in \mathcal{O}(n^2)$ denote the affine segments of φ such that, for each $j \in \{1, \dots, q\}$, the function φ takes on some affine form $\varphi(\lambda) = \alpha^{(j)} + \lambda \cdot \beta^{(j)}$ for $\lambda \in I_j$.

Now consider one such interval I_j . If $\varphi(\lambda) = 0$ for $\lambda \in I_j$, it also holds that $p^*(\lambda) = 0$ since $\varphi(\lambda) \geq \frac{1}{2} \cdot p^*(\lambda)$ for $\lambda \in \mathbb{R}$, so the all-zero solution is optimal. Otherwise, for the non-rounded scaled profit of each item i , it holds that

$$\frac{n \cdot p_i(\lambda)}{\varepsilon \cdot \varphi(\lambda)} = \frac{n \cdot (a_i + \lambda \cdot b_i)}{\varepsilon \cdot (\alpha^{(j)} + \lambda \cdot \beta^{(j)})} =: \frac{n}{\varepsilon} \cdot f_i(\lambda).$$

²Note that items with a negative profit will not be present in an optimal solution.

These functions f_i are monotonous, since the first derivative fulfills

$$\begin{aligned}\frac{df_i}{d\lambda}(\lambda) &= \frac{b_i \cdot (\alpha^{(j)} + \lambda \cdot \beta^{(j)}) - (a_i + \lambda \cdot b_i) \cdot \beta^{(j)}}{(\alpha^{(j)} + \lambda \cdot \beta^{(j)})^2} \\ &= \frac{b_i \cdot \alpha^{(j)} - a_i \cdot \beta^{(j)}}{(\alpha^{(j)} + \lambda \cdot \beta^{(j)})^2}\end{aligned}$$

and, thus, does not change its sign within I_j . Hence, within the interval I_j , each scaled profit \tilde{p}_i has a monotone behavior. Moreover, it holds that $0 \leq f_i(\lambda) \leq 2$ for all $\lambda \in I_j$ since each item either has a non-negative profit within the whole interval or it will be ignored and since $p_i(\lambda) \leq p^*(\lambda) \leq 2 \cdot \varphi(\lambda)$. These observations yield that each scaled profit \tilde{p}_i changes its (integral) value at most $\mathcal{O}(\frac{n}{\epsilon})$ times within I_j since we only need to consider values for \tilde{p}_i between 0 and $\frac{2n}{\epsilon}$. Hence, the above recursive formulae only change $\mathcal{O}(\frac{n^2}{\epsilon^2})$ times within each I_j , in which case we have to repeat the computation of the values $w(i, p)$. This yields a total computational overhead of $\mathcal{O}(\frac{n^4}{\epsilon^2})$ per interval and, since there are at most $\mathcal{O}(n^2)$ intervals, a total running time of $\mathcal{O}(\frac{n^6}{\epsilon^2})$ for the parametric FPTAS. This running time will be significantly improved in the next subsection.

It should be noted that we need to take care of a proper definition of the returned intervals: For example, consider two scaled profits of the forms $\tilde{p}_i = \lfloor 1 + \lambda \rfloor$ and $\tilde{p}_j = \lfloor 1 - \lambda \rfloor$. For the critical value $\lambda_1 = 1$, both profits evaluate to 1. However, for $\lambda'_1 := \lambda_1 + \delta$ and $\lambda''_1 := \lambda_1 - \delta$ for a small value of δ , one of the profits already changes its integral value and the dynamic programming scheme may behave differently. One simple solution is to assess that we add a single-point interval $[\lambda_1, \lambda_1]$ for each critical value as well as two open intervals of the forms (λ_0, λ_1) and (λ_1, λ_2) , where λ_0 and λ_2 are adjacent critical values. The returned (ordered) sequence of intervals then alternates between single-point intervals and open intervals. For an open interval, we can obtain an approximate solution by setting λ to the middle point of the interval and performing the dynamic program for the corresponding constant scaled profits.

3.3. Improved Analysis

The major drawback of the above algorithm is that we basically need to reset the whole procedure whenever the function φ changes its behavior. With this approach, we were able guarantee that each scaled profit has a monotone behavior such

that each possible integral value is only attained at most once per interval. As it will be shown in this Section, we are somewhat allowed to “ignore” these changes without losing the guarantee that each possible value of the scaled profits will only be attained a constant number of times.

Theorem 1. *For each item $i \in \{1, \dots, n\}$, the scaled profits $\tilde{p}_i(\lambda)$ attain each value in $\{0, \dots, \frac{2n}{\epsilon}\}$ at most three times as λ increases from $-\infty$ to $+\infty$.*

PROOF. In order to prove the claim, it suffices to show that the sign of the first derivative of each function f_i changes at most twice while λ increases. As above, let I_1, \dots, I_q denote the intervals for which φ takes on some affine form $\varphi(\lambda) = \alpha^{(j)} + \lambda \cdot \beta^{(j)}$ such that

$$f_i(\lambda) = \frac{a_i + \lambda \cdot b_i}{\alpha^{(j)} + \lambda \cdot \beta^{(j)}}$$

for $\lambda \in I_j$. Since φ is convex and continuous, it holds that $\beta^{(j)} \leq \beta^{(j+1)}$ for $j \in \{1, \dots, q-1\}$ and that there is some index h such that $\alpha^{(j)} \leq \alpha^{(j+1)}$ for $j \in \{1, \dots, h\}$ and $\alpha^{(j)} \geq \alpha^{(j+1)}$ for $j \in \{h+1, \dots, q-1\}$. In fact, due to the construction of the intervals, these inequalities hold in the strict sense since $\beta^{(j)} = \beta^{(j+1)}$ would also imply that $\alpha^{(j)} = \alpha^{(j+1)}$ by continuity of φ , so both segments would belong to the same interval. Hence, if we plot the points $(\beta^{(j)}, \alpha^{(j)})^T$ into a β - α -space, we get a picture as shown in Figure 2. Moreover, for each $j \in \{1, \dots, q-1\}$, there is some $\lambda_j \in \mathbb{R}$ with

$$\alpha^{(j)} + \lambda_j \cdot \beta^{(j)} = \alpha^{(j+1)} + \lambda_j \cdot \beta^{(j+1)}$$

due to the continuity and construction of φ . Hence, since $\beta^{(j+1)} = \beta^{(j)} + \delta_j$ for some value $\delta_j > 0$, we get that

$$\begin{aligned}\alpha^{(j+1)} &= \alpha^{(j)} + \lambda_j \cdot \beta^{(j)} - \lambda_j \cdot \beta^{(j+1)} \\ &= \alpha^{(j)} - \lambda_j \cdot \delta_j,\end{aligned}$$

so the slope of the line that connects the points $(\beta^{(j)}, \alpha^{(j)})^T$ and $(\beta^{(j+1)}, \alpha^{(j+1)})^T$ evaluates to

$$\frac{\alpha^{(j+1)} - \alpha^{(j)}}{\beta^{(j+1)} - \beta^{(j)}} = \frac{-\lambda_j \cdot \delta_j}{\delta_j} = -\lambda_j$$

and, thus, decreases while j increases. This yields that the piecewise linear function g connecting each of the points $(\beta^{(j)}, \alpha^{(j)})^T$ in the order $j = 1, \dots, q$

is concave³ (as illustrated by the highlighted area in Figure 2).

Now, for some specific item $i \in \{1, \dots, n\}$, consider the first derivate of f_i , which as we have seen evaluates to

$$\frac{df_i}{d\lambda}(\lambda) = \frac{b_i \cdot \alpha^{(j)} - a_i \cdot \beta^{(j)}}{(\alpha^{(j)} + \lambda \cdot \beta^{(j)})^2}$$

as shown above. Since the denominator is always positive, we need to bound the number of times the sign of the numerator changes. The value $b_i \cdot \alpha^{(j)} - a_i \cdot \beta^{(j)}$ can be interpreted as the inner product of the vectors $(-a_i, b_i)$ and $(\beta^{(j)}, \alpha^{(j)})^T$. Hence, since $(-a_i, b_i) \cdot (b_i, a_i)^T = 0$, the sign of the derivate changes whenever the function g crosses the line going through the origin and the point $(b_i, a_i)^T$ (see the dotted line in Figure 2). Since g is concave as shown above, this can happen at most two times while λ increases, which yields the claim. \square

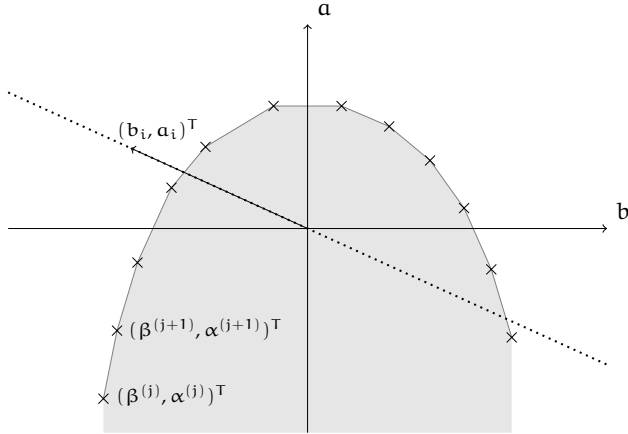


Figure 2: Plotting the slope $\beta^{(j)}$ and intersect $\alpha^{(j)}$ of each affine segment of φ . The piecewise linear function connecting these points is concave and intersects with each straight line through the origin at most twice.

Theorem 1 shows that each possible profit is only attained a constant number of times per item *al-*

³This can also be seen by arguments used in the field of computational geometry: It is known that the upper envelope of a set of affine functions of the form $c \cdot \lambda - d$ corresponds to the lower surface of a convex hull in the *dual space*, which is clearly convex. Such a dual space contains a point $(c, d)^T$ for each affine function of the above form in the primal space and, conversely, an affine function $\lambda \cdot c - \mu$ for each point $(\lambda, \mu)^T$ in the primal space. Hence, each line segment (breakpoint) of our upper envelope corresponds to a corner point (line segment) of the lower surface of a convex hull in the dual space (cf. [22]). In fact, Figure 2 shows the dual space mirrored at the b -axis.

though the involved functions f_i are rational functions whose denominator changes for increasing λ . Hence, each item only creates $\mathcal{O}(\frac{n}{\epsilon})$ subintervals as opposed to the $\mathcal{O}(n^2 \cdot \frac{n}{\epsilon})$ subintervals proven in Section 3.2. It remains to show that we can determine these subintervals efficiently.

As shown in the proof of Theorem 1, the slope of each function f_i changes at most twice, yielding for each item up to three partitions of the set of intervals I_1, \dots, I_q such that f_i is monotonous within each partition. By scanning through the sequence of intervals of φ , we can determine these three partitions for all items in total time $\mathcal{O}(n \cdot n^2) = \mathcal{O}(n^3)$. For each item, each partition, and each possible scaled profit in $\{0, \dots, \frac{2n}{\epsilon}\}$ (which can be attained only once in the partition), we perform a binary search on the intervals in the partition in order to find a value for λ at which the corresponding profit is attained, if such a λ exists. This can be done in $\mathcal{O}(n \cdot 3 \cdot \frac{n}{\epsilon} \cdot \log n^2) = \mathcal{O}(\frac{n^2}{\epsilon} \cdot \log n)$ time in total. Finally, we need to sort this list of critical values of λ in order to determine the subintervals of the FPTAS, which can be done in $\mathcal{O}(n \cdot \frac{n}{\epsilon} \cdot \log(n \cdot \frac{n}{\epsilon})) = \mathcal{O}(\frac{n^2}{\epsilon} \cdot \log \frac{n}{\epsilon})$ time.

In summary, we need $\mathcal{O}(n^3 + \frac{n^2}{\epsilon} \cdot \log \frac{n}{\epsilon})$ time to determine the $\mathcal{O}(\frac{n^2}{\epsilon})$ subintervals of the FPTAS. For each of these subintervals, we need to perform the traditional FPTAS with the corresponding (constant) scaled profits, which can be done in $\mathcal{O}(\frac{n^2}{\epsilon})$ time each. Hence, we obtain an FPTAS for the parametric knapsack problem running in $\mathcal{O}(\frac{n^4}{\epsilon^2})$ time in total. This yields the main result of this paper:

Theorem 2. *There is an FPTAS for the parametric knapsack problem running in strongly polynomial time $\mathcal{O}(\frac{n^4}{\epsilon^2})$.* \square

4. Combining FPTASs

In the previous section, we have seen that the parametric knapsack problem can be divided into $\mathcal{O}(\frac{n^2}{\epsilon})$ subproblems, for which we need to provide $(1 - \epsilon)$ -approximate solutions. These subproblems were created in a way such that the scaled profits are constant for each subproblem. Each of them can be seen as a new, independent, and non-parametric knapsack instance (albeit a special one, since the profits are now of polynomial size). In Section 3, we simply solved each of the resulting knapsack instances exactly in $\mathcal{O}(\frac{n^2}{\epsilon})$ time. The main observation

of this section is that we actually do not necessarily need to solve the subproblems exactly – it suffices to solve them up to a factor of $(1 - \varepsilon)$ using *any* FPTAS for the traditional knapsack problem.

Consider one fixed interval I' of the $\mathcal{O}(\frac{n^2}{\varepsilon})$ subintervals of the problem. For each $\lambda \in I'$, the scaled profits \tilde{p}_i take on constant values. Moreover, it holds that $\frac{1}{2} \cdot p^*(\lambda) \leq \varphi(\lambda) \leq p^*(\lambda)$ for any $\lambda \in I'$ as shown in Section 2.2. Let \bar{x} denote a solution returned by some FPTAS for the traditional knapsack problem that is called on an instance with the scaled profits and let x denote an exact solution to the scaled instance (which, e.g., can be obtained by the dynamic programming scheme as above). Clearly, it holds that

$$\tilde{p}(\bar{x}) := \sum_{i=1}^n \tilde{p}_i \cdot \bar{x}_i \geq (1 - \varepsilon) \cdot \sum_{i=1}^n \tilde{p}_i \cdot x_i =: \tilde{p}(x).$$

For any fixed $\lambda \in I'$ and an optimal solution x^* for the unscaled problem at λ , we then get the following approximation guarantee for the solution \bar{x} :

$$\begin{aligned} p(\bar{x}) &= \sum_{i=1}^n p_i \cdot \bar{x}_i \geq \sum_{i=1}^n M \cdot \left\lfloor \frac{p_i}{M} \right\rfloor \cdot \bar{x}_i \\ &= M \cdot \tilde{p}(\bar{x}) \\ &\geq (1 - \varepsilon) \cdot M \cdot \tilde{p}(x) \\ &\geq (1 - \varepsilon) \cdot M \cdot \tilde{p}(x^*) \\ &\geq (1 - \varepsilon) \cdot M \cdot \sum_{i=1}^n \left(\frac{p_i}{M} - 1 \right) \cdot x_i^* \\ &= (1 - \varepsilon) \cdot \left(p^*(\lambda) - M \cdot \sum_{i=1}^n x_i^* \right) \\ &\geq (1 - \varepsilon) \cdot (p^*(\lambda) - \varepsilon \cdot \varphi(\lambda)) \\ &\geq (1 - \varepsilon) \cdot (p^*(\lambda) - \varepsilon \cdot p^*(\lambda)) \\ &\geq (1 - \varepsilon)^2 \cdot p^*(\lambda) \\ &= (1 - 2\varepsilon + \varepsilon^2) \cdot p^*(\lambda) \\ &\geq (1 - 2\varepsilon) \cdot p^*(\lambda). \end{aligned}$$

Setting $\varepsilon' := \frac{\varepsilon}{2}$ then yields the desired approximation guarantee. Hence, although the subproblems were designed in a way such that the basic dynamic programming scheme does not change its behavior, we do not necessarily need to execute it but can also use an FPTAS instead.

Theorem 3. *There is an FPTAS for the parametric knapsack problem running in $\mathcal{O}(\frac{n^2}{\varepsilon} \cdot A(n, \varepsilon))$ time, where*

$A(n, \varepsilon)$ denotes the running time of an FPTAS for the traditional knapsack problem. \square

Note that it clearly holds that $A(n, \varepsilon) \in \Omega(n)$, so the running time of the main procedure will dominate the overheads to compute φ and the set of subintervals.

At present, the best FPTAS for the traditional knapsack problem is given by Kellerer and Pferschy [4, 5] and achieves a running time of

$$\mathcal{O} \left(n \cdot \min \left\{ \log n, \log \frac{1}{\varepsilon} \right\} + \frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon} \cdot \min \left\{ n, \frac{1}{\varepsilon} \log \frac{1}{\varepsilon} \right\} \right).$$

Under the commonly used assumption that n is much larger than $\frac{1}{\varepsilon}$ in practice [21], this running time evaluates to

$$\mathcal{O} \left(n \log \frac{1}{\varepsilon} + \frac{1}{\varepsilon^3} \log^2 \frac{1}{\varepsilon} \right),$$

yielding an FPTAS for the parametric knapsack problem with a strongly polynomial running time of

$$\mathcal{O} \left(\frac{n^3}{\varepsilon} \log \frac{1}{\varepsilon} + \frac{n^2}{\varepsilon^4} \log^2 \frac{1}{\varepsilon} \right).$$

References

- [1] O. Ibarra, C. Kim, Fast approximation algorithms for the knapsack and sum of subset problems, *Journal of the ACM (JACM)* 22 (4) (1975) 463–468.
- [2] E. L. Lawler, *Combinatorial optimization: networks and matroids*, Courier Corporation, 2001.
- [3] M. Magazine, O. Ogun, A fully polynomial approximation algorithm for the 0–1 knapsack problem, *European Journal of Operational Research* 8 (3) (1981) 270–273.
- [4] H. Kellerer, U. Pferschy, A new fully polynomial time approximation scheme for the knapsack problem, *Journal of Combinatorial Optimization* 3 (1) (1999) 59–71.
- [5] H. Kellerer, U. Pferschy, Improved dynamic programming in connection with an FPTAS for the knapsack problem, *Journal of Combinatorial Optimization* 8 (1) (2004) 5–11.
- [6] H. Kellerer, U. Pferschy, D. Pisinger, *Knapsack Problems*, Springer, 2004.
- [7] P. Carstensen, Complexity of some parametric integer and network programming problems, *Mathematical Programming* 26 (1) (1983) 64–75.
- [8] R. Karp, J. Orlin, Parametric shortest path algorithms with an application to cyclic staffing, *Discrete Applied Mathematics* 3 (1) (1981) 37–45.
- [9] N. Young, R. Tarjan, J. Orlin, Faster parametric shortest path and minimum-balance algorithms, *Networks* 21 (2) (1991) 205–221.

- [10] P. Carstensen, The complexity of some problems in parametric linear and combinatorial programming .
- [11] K. Mulmuley, P. Shah, A lower bound for the shortest path problem, in: Computational Complexity, IEEE, 14–21, 2000.
- [12] D. Fernández-Baca, G. Slutzki, D. Eppstein, Using sparsification for parametric minimum spanning tree problems, in: Scandinavian Workshop on Algorithm Theory, Springer, 149–160, 1996.
- [13] P. Agarwal, D. Eppstein, L. Guibas, M. Henzinger, Parametric and kinetic minimum spanning trees, in: Foundations of Computer Science, 1998. Proceedings. 39th Annual Symposium on, IEEE, 596–605, 1998.
- [14] G. Gallo, M. Grigoriadis, R. Tarjan, A fast parametric maximum flow algorithm and applications, SIAM Journal on Computing 18 (1) (1989) 30–55.
- [15] S. McCormick, Fast algorithms for parametric scheduling come from extensions to parametric maximum flow, Operations Research 47 (5) (1999) 744–756.
- [16] M. Scutella, A note on the parametric maximum flow problem and some related reoptimization issues, Annals of Operations Research 150 (1) (2007) 231–244.
- [17] A. Giudici, P. Halffmann, S. Ruzika, C. Thielen, Approximation schemes for the parametric knapsack problem, Information Processing Letters 120 (2017) 11–15.
- [18] M. Eben-Chaime, Parametric solution for linear bicriteria knapsack models, Management Science 42 (11) (1996) 1565–1575.
- [19] T. Erlebach, H. Kellerer, U. Pferschy, Approximating multi-objective knapsack problems, in: Workshop on Algorithms and Data Structures, Springer, 210–221, 2001.
- [20] J. Hershberger, Finding the upper envelope of n line segments in $O(n \log n)$ time, Information Processing Letters 33 (4) (1989) 169–174.
- [21] E. Lawler, Fast approximation algorithms for knapsack problems, Mathematics of Operations Research 4 (4) (1979) 339–356.
- [22] M. D. Berg, M. V. Kreveld, M. Overmars, O. Schwarzkopf, Computational geometry algorithms and applications, Springer, 2000.